

---

The **Updates** on the Benchmarking Methodology for Intra-domain  
and Inter-domain Source Address Validation  
*draft-chen-bmwg-savnet-sav-benchmarking-08*

Li Chen, Dan Li, **Libin Liu**, Lancheng Qin

SAVNET WG Meeting, IETF 124

# Background

□ Defining the methodologies for benchmarking the performance of intra-domain and inter-domain source address validation (SAV) mechanism

□ Historical versions

◆ draft-chen-bmwg-savnet-sav-benchmarking-00, IETF 120 BMWG and SAVNET WG

◆ draft-chen-bmwg-savnet-sav-benchmarking-01, August 7, 2024

◆ draft-chen-bmwg-savnet-sav-benchmarking-02, IETF 121 BMWG and SAVNET WG

◆ draft-chen-bmwg-savnet-sav-benchmarking-03, IETF 122 BMWG

◆ draft-chen-bmwg-savnet-sav-benchmarking-04, May 20, 2025

◆ draft-chen-bmwg-savnet-sav-benchmarking-05, IETF 123 BMWG and SAVNET WG

◆ draft-chen-bmwg-savnet-sav-benchmarking-06, August 29, 2025

◆ draft-chen-bmwg-savnet-sav-benchmarking-07, October 20, 2025

◆ **draft-chen-bmwg-savnet-sav-benchmarking-08, IETF 124 BMWG and SAVNET WG**

# Overview of the Document

## □ SAV performance indicators

- ◆ False positive rate
- ◆ False negative rate
- ◆ Protocol convergence time
- ◆ Protocol message processing throughput
- ◆ Data plane SAV table refreshing rate
- ◆ Data plane forwarding rate
- ◆ Resource utilization

## □ Test cases of intra-domain and inter-domain SAV

- ◆ False positive and false negative rates
- ◆ Control plane performance
- ◆ Data plane performance

## □ Reporting format

1. Introduction . . . . .	3
1.1. Goal and Scope . . . . .	3
1.2. Requirements Language . . . . .	4
2. Terminology . . . . .	4
3. Test Methodology . . . . .	4
3.1. Test Setup . . . . .	4
3.2. Network Topology and Device Configuration . . . . .	5
4. SAV Performance Indicators . . . . .	6
4.1. False Positive Rate . . . . .	6
4.2. False Negative Rate . . . . .	6
4.3. Protocol Convergence Time . . . . .	6
4.4. Protocol Message Processing Throughput . . . . .	6
4.5. Data Plane SAV Table Refreshing Rate . . . . .	6
4.6. Data Plane Forwarding Rate . . . . .	7
4.7. Resource Utilization . . . . .	7
5. Benchmarking Tests . . . . .	7
5.1. Intra-domain SAV . . . . .	7
5.1.1. False Positive and False Negative Rates . . . . .	7
5.1.2. Control Plane Performance . . . . .	15
5.1.3. Data Plane Performance . . . . .	17
5.2. Inter-domain SAV . . . . .	18
5.2.1. False Positive and False Negative Rates . . . . .	19
5.2.2. Control Plane Performance . . . . .	32
5.2.3. Data Plane Performance . . . . .	32
5.3. Resource Utilization . . . . .	32
6. Reporting Format . . . . .	33
7. IANA Considerations . . . . .	33
8. Security Considerations . . . . .	33
9. References . . . . .	33
9.1. Normative References . . . . .	33
9.2. Informative References . . . . .	34
Acknowledgements . . . . .	35
Authors' Addresses . . . . .	35

# Comments & Our Response

□ C1: Comments on the test setup using spoofed traffic with different types of source addresses

◆ The setting of the forged source address may have an impact on the test results. An invalid packet may forge an unused source address, private network source address, internal source address, or external source address. Different SAV mechanisms vary in their ability to block packets with forged addresses of various types. (Nan Geng)

➤ **Our response:** In Sections 5.1.1 and 5.2.1, we have added the test setup which requires that under each scenario, the generated spoofed traffic SHOULD include different types of forged source addresses and the ratios among these different types of forged source addresses SHOULD vary.

# Comments & Our Response

□ C2: Comments on the test setup using different ratios of spoofed traffic to legitimate traffic

◆ For figure 5, how can we understand "The ratio of spoofed to legitimate traffic"? The ratio seems cannot directly reflect the protective effectiveness of SAV. (Nan Geng, Xueyan Song)

➤ **Our response:** In the mailing list, we have explained that the ratio of spoofed to legitimate traffic indicates the proportions of spoofed and legitimate traffic among the overall testing traffic. **We use this because the proportions of spoofed traffic may vary in practice, and we try to simulate the possible ratios of spoofed traffic. We have revised all the corresponding parts in the document and made tester generate varying ratios of spoofed traffic.** For example, in Figure 5, the testing traffic can generate spoofed traffic which simulates different ratios of source addresses belonging to the sub network.

# Comments & Our Response

## □ C3: Comments on the SAV performance indicators

◆ Section 4.6 Data Plane Forwarding Rate: Would the proportion of the decrease in forwarding rate (compared to disabling SAV in dataplane) be more appropriate? (Nan Geng)

➤ **Our response:** In Section 4.6, we have clarified the definition of data plane forwarding rate and added the suggestion to measure the data plane forwarding rates with enabling and disabling SAV, and evaluate their differences.

◆ The listed KPIs are important foundational metrics. But in real networks, some additional dimensions also matter. For example: stability over time (do KPIs fluctuate under sustained load?) and sensitivity to table size (does performance degrade as the SAV table grows large?). Adding such items as optional KPIs may better reflect real-world behavior. (Yuanyuan Zhang)

➤ **Our response:** In Section 4, we have added considerations to analyze the standard deviation of KPIs and measure SAV table refreshing rate and data plane forwarding rate with varying SAV table sizes.

# Comments & Our Response

## □ C4: Comments on testing control plane performance

◆ Not all SAV mechanisms actually have a dedicated control-plane protocol. In such cases, the "control-plane benchmarking" may not be applicable. I suggest making these control-plane tests optional, with wording like "only applicable if the SAV mechanism includes an explicit control-plane protocol." More generally, it would help if the draft explicitly distinguishes between mandatory tests (that all SAV mechanisms should support) and optional tests (only applicable in certain designs). (Yuanyuan Zhang)

➤ **Our response:** In sections 5.1.2 and 5.2.2, we have added the explanations that **the tests for control plane performance of the DUT which performs intra-domain and inter-domain SAV are OPTIONAL**. Only DUT which implements the SAV mechanism using an explicit control-plane communication protocol **SHOULD** be tested on its control plane performance.

# Comments & Our Response

## □C5: Comments on the figures and their descriptions

- ◆Figure 2 in Section 5.1.1: The route of 10.2.0.0/15 announced by whom is not mentioned. (Nan Geng)
- ◆Figure 7 in Section 5.1.1: It shows a 10.0.0.0/15 subnet, which actually covers both /16s. In this case, both prefixes should be legitimate, so the example is confusing. (Yuanyuan Zhang)
  - Our response:** In Section 5.1.1, we have updated the descriptions about Figures 2 and 7, as well as other similar parts in Section 5.1.

## □C6: Comments on clarifying some terminologies

- ◆The definitions of host-facing router and customer-facing router are not very clear. For host-facing router: maybe just say "an edge router directly connected to end-host subnets". For customer-facing router: do you mean the classic CE router at the AS boundary? (Yuanyuan Zhang)
  - Our response:** We have revised the definitions of host-facing router and customer-facing router in Section 2, which are consistent with the intra-domain SAVNET PS document. Host-facing router is an edge router directly connected to a layer-2 host network, and customer-facing router is an edge router connected to a non-BGP customer network which includes routers and runs the routing protocol.

# Comments & Our Response

## □C7: Comments on adding more real-world testing scenarios

- ◆ Would FRR be considered in the future? The FRR scenario can also be considered for testing the performance of SAV. (Nan Geng)
- ◆ Should some common real-world cases be added? For example: Complicated routing policies (FRR, policy-based routing, community rewriting). (Yuanyuan Zhang)
  - **Our response:** In Sections 5.1 and 5.2, we have added four new testing scenarios including intra-domain SAV under FRR scenario, intra-domain SAV under PBR scenario, inter-domain SAV under FRR scenario, and inter-domain SAV under PBR scenario.

## □C8: Comments on the partial deployment scenarios

- ◆ The partial deployment scenarios, where only a subset of routers or ASes implement SAV, should be considered. (Yuanyuan Zhang)
  - **Our response:** All the scenarios in this document are partial deployment, i.e., enabling SAV with a subset of routers for intra-domain network or partial ASes for inter-domain network.

# Summary of Main Updates (v-08 vs. v-05)

## □ Terminology

- ◆ Updating the definitions of *host-facing router* and *customer-facing router*

## □ SAV performance indicator

- ◆ Data plane SAV table refreshing rate
- ◆ Data plane forwarding rate

## □ Test cases of intra-domain and inter-domain SAV

- ◆ Adding different types of spoofed source addresses
- ◆ Illustrating different ratios of spoofed to legitimate traffic
- ◆ Illustrating control plane performance testing
- ◆ Adding new testing cases for SAV under FRR and PBR scenarios

1. Introduction . . . . .	3
1.1. Goal and Scope . . . . .	3
1.2. Requirements Language . . . . .	4
2. Terminology . . . . .	4
3. Test Methodology . . . . .	4
3.1. Test Setup . . . . .	4
3.2. Network Topology and Device Configuration . . . . .	5
4. SAV Performance Indicators . . . . .	6
4.1. False Positive Rate . . . . .	6
4.2. False Negative Rate . . . . .	6
4.3. Protocol Convergence Time . . . . .	6
4.4. Protocol Message Processing Throughput . . . . .	6
4.5. Data Plane SAV Table Refreshing Rate . . . . .	6
4.6. Data Plane Forwarding Rate . . . . .	7
4.7. Resource Utilization . . . . .	7
5. Benchmarking Tests . . . . .	7
5.1. Intra-domain SAV . . . . .	7
5.1.1. False Positive and False Negative Rates . . . . .	7
5.1.2. Control Plane Performance . . . . .	15
5.1.3. Data Plane Performance . . . . .	17
5.2. Inter-domain SAV . . . . .	18
5.2.1. False Positive and False Negative Rates . . . . .	19
5.2.2. Control Plane Performance . . . . .	32
5.2.3. Data Plane Performance . . . . .	32
5.3. Resource Utilization . . . . .	32
6. Reporting Format . . . . .	33
7. IANA Considerations . . . . .	33
8. Security Considerations . . . . .	33
9. References . . . . .	33
9.1. Normative References . . . . .	33
9.2. Informative References . . . . .	34
Acknowledgements . . . . .	35
Authors' Addresses . . . . .	35

---

Thanks!